AD-A133 657 ORDERING METHODS FOR SPARSE MATRICES(U) BOEING COMPUTER SERVICES CO TUKWILA WA ENERGY TECHNOLOGY APPLICATIONS DIV J G LEWIS 28 JUL 83 AFOSR-TR-83-0819 1/1

UNCLASSIFIED F49620-81-C-0072 F/G 12/1 NL

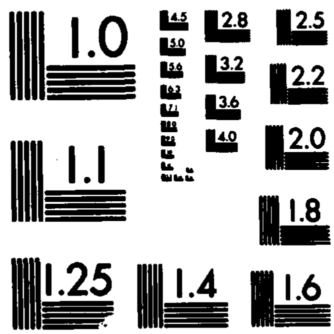

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AFOSR-TR- 83-0819

A133 657

(9)

ORDERING METHODS FOR SPARSE MATRICES: FINAL REPORT

AFOSR Contract F49620-81-C-0072

Report No. 2

Final Report

by

John G. Lewis

Boeing Computer Services Company

July 28, 1983

KEY WORDS AND PHRASES: linear equations, sparse matrices, reordering algorithms, Hellerman-Rarick algorithm, P4 algorithm, unsymmetric matrices, Gauss elimination.

Document describes the performance of

ABSTRACT: This report summarizes the activities at Boeing Computer Services Company on AFOSR Contract F49620-81-C-0072 from July 1, 1981 until July 28, 1983. Five tasks were performed: creation of a comprehensive test matrix collection, analysis of the Hellerman-Rarick P4 algorithm, production of a P4 code, production of a diagnostic code, and comparative analysis of several algorithms using the test matrices and the diagnostic code. Reports on the completion of the five tasks are given, relevant reports and publications of project personnel are listed and related sparse matrix activities are discussed.

INTRODUCTION

The solution of systems of large sparse linear equations is a fundamental computational step in the numerical solution of many scientific and engineering problems. These problems arise in such diverse areas as structural analysis, linear programming, network analysis, chemical process design, electromagnetic pulse (EMP) analysis, optimization, steady state analysis, and policy planning. When direct solution methods are used to solve these equations one of the major difficulties is choosing a reordering of the rows and columns of the sparse matrix to reduce some measure of solution cost.

Because sparse matrix research has grown independently out of many disciplines, there are many different heuristic methods (band, profile, Markowitz, tearing, P4, and variations) presently used to accomplish this reordering. The challenge in building standard software is to determine if one heuristic works adequately across a broad class of problems, or if several heuristics must be available in a general purpose code. If several heuristics are needed, matrix classes must be identified as a basis for matching a given matrix with the proper ordering method. In either case it also must be determined how much performance will improve for particular problem classes if specialized sparse matrix code rather than general purpose code is used.

This research project was concerned with answering these questions. This report describes the current status of the project. The following topics are covered by this report: research objectives, status of the research effort, relevant publications by the project personnel, professional personnel associated with the research effort, and related sparse matrix activities at Boeing Computer Services.

RESEARCH OBJECTIVES

The research objectives were broken down into five tasks. The tasks, as described in the future tense at the beginning of the research project, were:

<u>Task 1: Creation of a Comprehensive Test Matrix Collection.</u> The goal of this task is to assemble a large collection of sparse matrices. These matrices will be representative of realistic problems arising in many different application areas. They should be of varying sizes and structural characteristics. Considerable effort is going into this task in order to have a comprehensive collection.

<u>Task 2: Analyze the Hellerman-Rarick P4 Algorithm.</u> In this task the Hellerman-Rarick reordering algorithm for sparse unsymmetric linear equations and all of its known variations will be studied with the intent of understanding the applicability, stability, and effectiveness of it. A precise algorithmic description of the most important version will be produced. This description will form the basis for studying the different variations.

Task 3: Producing a P4 Code. The precise algorithmic description of the Hellerman-Rarick algorithm from Task 2 will be used to produce a high quality FORTRAN code implementing the algorithm. This code will include all modifications which are improvements of the original algorithm.

Task 4: Producing a Diagnostic Code. A diagnostic program will be produced which will be capable of monitoring the effectiveness of various ordering algorithms applied to sparse matrices. The diagnostic code will monitor many characteristics such as accuracy, fill, storage, run-time and operation counts.

Task 5: Comparative Analysis. This final task is to utilize the test collection (Task 1) and the diagnostic program (Task 4) to provide detailed comparisons of various ordering algorithms including the Hellerman-Rarick algorithm (Task 3) and the MA28 code from the Harwell library.

STATUS OF THE RESEARCH EFFORT

All of the tasks were completed as planned. Tasks 2 and 4 were completed during the first year of this project. Task 1 is open-ended; work on Task 1 was performed throughout the contract, and will continue under BCS auspices indefinitely. Task 3 and 5 were completed in the second year of the project. This section describes those activities.

Task 1. We have contacted many people from different application areas who have sparse matrix problems. Test problems from structural engineering, chemical process design, linear programming, circuit analysis and optimization have been received. We published a request for test matrices in the ACM SIGNUM Newsletter and in the IMA IMANA Newsletter. We actively recruited test problems at recent professional meetings including the SIAM Conference on Applied Linear Algebra, the SIAM 30th Anniversary Meeting and the Sparse Matrix Symposium. This task continues to be performed in collaboration with Dr. Iain Duff of AERE Harwell in England.

Task 2. We have analyzed the Hellerman-Rarick P4 algorithm and have produced a very simple algorithmic description of it. Previous descriptions have been rather lengthy and complicated and, thus, have made analysis difficult. Our simplified description has led to an understanding of how the algorithm works and what types of problems cause it difficulties. Because of the simplified description, we also have been able to analyze modifications of the algorithm. Several modifications seem very promising. In fact, they cannot break down because of intermediate structural singularity, a problem with other versions of the Hellerman-Rarick algorithm. This work was discussed at the SIAM Meeting at Stanford University, July 19-23, 1982, and the Sparse Matrix Symposium at Fairfield Glade, October 25-27, 1982. These modifications were tested for their effectiveness on the test matrix collection in Task 5. The modified algorithm performed extremely well on one measure of effectiveness, but not on another. The algorithmic description and a synopsis of the evaluation were submitted to SIAM Journal of Numerical Analysis. The evaluations were more fully discussed as part of Task 5.

Task 3. We evaluated two different existing Hellerman-Rarick P4 codes: one from Mike Saunders of Stanford University and one from the World Bank. We have converted both of these codes so that they can interact with our diagnostic code (task 4). The World Bank code was modified in accordance with the results of Task 2, and both the original World Bank ordering and our modification were used in Task 5.

Task 4. The diagnostic program was designed, written, and tested. It collects and analyzes data concerning the performance of the different ordering algorithms. This information is instrumental in deciding which algorithms are most effective for various matrix problems in our collection of test matrices. The types of data gathered include: order of the matrix, number of nonzeros in different parts of the matrix, nearness to symmetry, measures of regularity of distribution of nonzeros, storage required, fill, execution times, operation counts, stability bounds and relative density. The diagnostic program measures the ability of ordering algorithms to reduce the amount of computer memory required to store the factored matrix as well as reducing the number of operations for the actual factorization. This program was used to monitor the results of the algorithms tested in task 5. Copies of this program will be sent to other sparse matrix researchers.

Task 5. Fifty-six test matrices were used as a test set to evaluate the performance of the Markowitz (MA28) ordering, the Hellerman-Rarick P^4 ordering and our modification of the latter, which we call the P^5 ordering. The results of this task were summarized in the paper "Evaluation of Orderings for Unsymmetric Sparse Matrices", which will be submitted to the SIAM Journal of Scientific and Statistic Computing. The results indicated that the modified Hellerman-Rarick algorithm developed under this contract was superior in the reduction of fill, but that the Markowitz algorithm was superior in reducing fill while preserving numerical stability. Interesting results on the dependence of ordering effectiveness on matrix properties were obtained. The massive amounts of data made analysis difficult. A novel approach, statistical analysis, proved to be very useful in summarizing and interpreting the data, and provided insight into the properties of sparse matrices.

RECENT RELEVANT REPORTS AND PUBLICATIONS OF THE PROJECT PERSONNEL

Benjamin Dembart, David S. Dodson, John G. Lewis, Kenneth W. Neves and William G. Poole, Jr., "Vector Computing: Requirements and Constraints", Proceedings of the Eighth Conference on Electronic Computation, J. K. Nelson, Jr,. Ed., Houston, Tx., 1983, pp. 28-36.

David S. Dodson, John G. Lewis and William G. Poole, Jr., "Tailoring Mathematical Software for the CRAY-1", Proceedings of the Conference on the Computing Environment for Mathematical Software, Pasadena, California, 1981, pp. 27-28.

David S. Dodson and John G. Lewis, "Improving the Performance of a Sparse Matrix Solver on the CRAY-1", Proceedings of the Science, Engineering and the CRAY-1 Conference, Minneapolis, Mn., 1982.

David S. Dodson and Roger G. Grimes, "Remark on Algorithm 539," ACM Transactions on Mathematical Software 8, 4 (1982), pp. 403-404.

Iain S. Duff, Roger G. Grimes, John G. Lewis and William G. Poole, Jr., "Sparse Matrix Test Problems," ACM SIGNUM Newsletter, Volume 17, No. 2, June 1982, page 22, and IMANA Newsletter, Volume 6, No. 3, April 1982, page 19.

A. M. Erisman, R. G. Grimes, J. G. Lewis and W. G. Poole, Jr., "A Structurally Stable Modification of Hellerman-Rarick's P^4 Algorithm for Reordering Unsymmetric Sparse Matrices," BCS Mathematics and Modeling Technical Report MM-3, April 1983, submitted to SIAM J. of Numer. Anal.

A. M. Erisman, R. G. Grimes, J. G. Lewis, W. G. Poole, Jr. and H. D. Simon, "Evaluation of Orderings for Unsymmetric Sparse Matrices," BCS Mathematics and Modeling Technical Report MM-5, July 1983, submitted to SIAM J. Sci. Stat. Comp.

A. M. Erisman, "Sparse Matrix Problems in Electric Power System Analysis," in Sparse Matrices and their Uses, Iain S. Duff, editor, Academic Press, New York, 1981.

Roger G. Grimes and John G. Lewis, "Condition Number Estimation for Sparse Matrices," SIAM Journal on Scientific and Statistical Computing, Volume 2, No. 4, December 1981, pages 384-388.

John G. Lewis and Roger G. Grimes, "Practical Lanczos Algorithms for Structural Engineering Eigenvalue Problems", in Sparse Matrices and their Uses, Iain S. Duff, editor, Academic Press, New York, 1981.

John G. Lewis, "ALGORITHM 582: The Gibbs-Poole-Stockmeyer and Gibbs-King Algorithms for Reordering Sparse Matrices," ACM Transactions on Mathematical Software, Vol. 8, No. 2, June 1982, pages 190-194.

John G. Lewis, "Implementation of the Gibbs-Poole-Stockmeyer and Gibbs-King Algorithms," ACM Transactions on Mathematical Software, Vol. 8, No. 2, June 1982, pages 180-189.

John G. Lewis, "Remark on Algorithms 508 and 509," ACM Transactions on Mathematical Software, Vol. 8, No. 2, June 1982, page 221.

John G. Lewis and Ronald G. Rehm, "The Numerical Solution of a Pressure Equation by Hybrid Conjugate Gradients", Journal of Research, National Bureau of Standards 85, 5, pp. 367-390.

John G. Lewis, "Numerical Experiments with SPARSPAK", to appear in SIGNUM Newsletter 18, 3 (1983).

John G. Lewis and William G. Poole, Jr., "Ordering Algorithms Applied to Sparse Matrices in Electric Power Problems," in Electric Power Problems: The Mathematical Challenge, A. M. Erisman, K. W. Neves and M. H. Dwarakanath, editors, SIAM, Philadelphia, 1980.

Horst D. Simon (with B. Parlett and L. Stringer), "Estimating the Largest Eigenvalue of a Symmetric Positive Definite Matrix with the Lanczos Algorithm," Math. of Comp. 38 (1983), pp. 153-165.

Horst D. Simon, "Analysis of the Symmetric Lanczos Algorithm with Reorthogonalization Methods,", submitted to Lin. Alg. Appl.

Horst D. Simon, "The Lanczos Algorithm with Partial Reorthogonalization," to appear in Math of Comp. (Jan. 1984).

Horst D. Simon, "Banded Preconditioning for the Solution of Symmetric Positive Definite Linear Systems," with B. Nour-Omid, Report, Dept. of Appl. Math, SUNY Stony Brook, 1982, submitted to SIAM J. Sci. Stat. Comp.

Elizabeth L. Yip, "A Sparse Capacitance Matrix Method," submitted to Math of Comp.

RECENT PRESENTATIONS AT PROFESSIONAL MEETINGS

Iain S. Duff, Roger G. Grimes, John G. Lewis, and William G. Poole, Jr., "Sparse Matrix Test Problems," presented at the Sparse Matrix Symposium 1982, Fairfield Glade, Tennessee, October 25-27, 1982.

Albert M. Erisman, "Matrix Modification and Partitioning," presented at the Sparse Matrix Symposium 1982, Fairfield Glade, Tennessee, October 25-27, 1982.

Albert M. Erisman, Roger G. Grimes, John G. Lewis, and William G. Poole, Jr., "The Hellerman-Rarick P4 Algorithm for Reordering Unsymmetric Sparse Matrices. Part I: The Algorithm," presented at the SIAM 30th Anniversary Meeting, Stanford University, Stanford, California, July 19-23, 1982.

Albert M. Erisman, Roger G. Grimes, John G. Lewis, and William G. Poole, Jr., "The Hellerman-Rarick P4 Algorithm for Reordering Unsymmetric Sparse Matrices. Part II: Structural Singularity and Modifications," presented at the SIAM 30th Anniversary Meeting, Stanford University, Stanford, California, July 19-23, 1982.

Albert M. Erisman, Roger G. Grimes, John G. Lewis, and William G. Poole, Jr., "A Structurally Stable Modification of the Hellerman-Rarick Algorithm for Reordering Unsymmetric Sparse Matrices," presented at the Sparse Matrix Symposium 1982, Fairfield Glade, Tennessee, October 25-27, 1982.

Roger G. Grimes, John G. Lewis, and William G. Poole, Jr., "Program for the Comparison of Reordering Algorithms for the Solution of Unsymmetric Sparse Systems of Equations," presented at the Sparse Matrix Symposium 1982, Fairfield Glade, Tennessee, October 25-27, 1982.

Albert M. Erisman, "Recent Developments in Sparse Matrix Computations with Application to Reservoir Simulation", presented at the BCS 1983 Spring Colloquium for the Geosciences, Dallas, May, 1983.

Kenneth W. Neves, "Vectorization of Scientific Software", presented at the NATO Workshop on High Speed Computation, Julich, W. Germany, June 1983.

Horst D. Simon, "The Lanczos Algorithm with Partial Reorthogonalization,", presented at the SIAM Conference on Applied Linear Algebra, Raleigh, April 1982.

Horst D. Simon, "The Lanczos Algorithm for the Solution of Nonsymmetric Linear Systems," presented at the Sparse Matrix Symposium 1982, Fairfield Glade, Tennessee, October 25-27, 1982.

Elizabeth L. Yip, "Preconditioning the Helmholtz Type Equation with the Fast Poisson Solver and the Capacitance Matrix," presented at the SIAM 1981 Fall Meeting, Cincinnati, October 1981.

Elizabeth L. Yip, "Conjugate Gradient Method for Unsymmetric Matrices Applied to Partial Differential Equations," presented at the SIAM 30th Anniversary Meeting, Palo Alto, June 1982.

RECENT ABSTRACTS PREPARED FOR PRESENTATION AT PROFESSIONAL MEETINGS

The following abstracts are preliminary reports of work in progress.

Kenneth W. Neves, "Vector Algorithms as a Function of Hardware Architecture," to be presented at the SIAM 1983 Conference on Parallel Processing for Scientific Computing, Norfolk, November 10-11, 1983.

John G. Lewis, "Vectorizing Direct Sparse Linear Equation Solvers," to be presented at the SIAM 1983 Conference on Parallel Processing for Scientific Computing, Norfolk, November 10-11, 1983.

PROFESSIONAL PERSONNEL ASSOCIATED WITH THE RESEARCH EFFORT

Three researchers from Boeing Computer Services, Albert M. Erisman, Roger G. Grimes and John G. Lewis performed most of the work over the two year period of this contract. William G. Poole, Jr., no longer with BCS, participated as program manager for the first year of the project. Horst D. Simon recently joined BCS and participated in the performance of task 5. Task 1, which is concerned with the creation of a comprehensive test matrix collection, is being carried out in collaboration with Dr. Iain S. Duff of AERE, Harwell, England. Kerry Whitaker of the ETA Applied Statistics group assisted with the statistical analyses in task 5.

RELATED SPARSE MATRIX ACTIVITIES AT BOEING COMPUTER SERVICES COMPANY

The mathematicians at Boeing Computer Services working on this project also are active in other projects which involve sparse matrix computations. This section briefly describes some of the most recent activities by those people. These projects are not funded by the AFOSR contract but they indicate the significant role that sparse matrix research plays at BCS.

Sparse Matrix Computations in Electric Power Problems. A. M. Erisman, R. G. Grimes, J. G. Lewis and W. G. Poole, Jr. have performed work recently which applied current sparse matrix technology to electric power problems. Two recently completed publications in this area are included in the earlier section on reports and publications. They appeared in the proceedings of two conferences: the SIAM meeting on Electric Power Problems: The Mathematical Challenge and the IMA conference on Sparse Matrices and their Uses.

Condition Number Estimation for Sparse Matrices. R. G. Grimes and J. G. Lewis have defined and implemented a condition number estimator for sparse matrices. The estimator has been implemented in a large, sparse eigenvalue program. A paper in the SIAM Journal on Scientific and Statistical Computing is mentioned in the section on publications.

Band and Envelope Reordering for Sparse Matrices. J. G. Lewis has greatly improved the Gibbs-Poole-Stockmeyer and Gibbs-King algorithm implementations for reducing the bandwidth and profile of a symmetric sparse matriix. Two papers and a short remark recently appeared in ACM Transactions on Mathematical Software (see section on publications).

Sparse Vector and Matrix Building Blocks for the CRAY-1. All five of the mathematicians on this project have been involved in developing, implementing and testing assembler language basic building blocks for sparse vector and matrix computations on the CRAY-1 computer. These subprograms are a part of CRAYPACK.

CRAY-1 Optimization of SPARSPAK and COMPLEX version of SPARSPAK. J. G. Lewis and R. G. Grimes have modified SPARSPAK so that a version optimized for the CRAY-1 and a COMPLEX version have been produced.

Problems from Structural Engineering. J. G. Lewis has tested several problems from structural engineering by exercising the various ordering algorithms in SPARSPAK. The preliminary results indicate that no single ordering algorithm will suffice for a general code. A Lanczos-based eigenvalue solver has also been developed for structural engineering codes. A version for the CRAY-1 computer has been prepared.

The Necessity of Pivoting. W. Kahan (of the University of California, Berkeley) and W. G. Poole, Jr. are working on a paper which discusses the conditions under which matrices do not need some form of pivoting for maintaining numerical stability.

SIAM Conference on Applied Linear Algebra. J. G. Lewis attended this conference in Raleigh, North Carolina in April 1982. At the conference, he publicized the sparse matrix collection and requested sample problems. H. D. Simon attended and presented a paper.

SIAM 30th Anniversary Meeting. A. M. Erisman, J. G. Lewis, W. G. Poole, Jr., H. D. Simon and E. L. Yip attended. Lewis, Poole and Yip presented papers. Several new examples for the sparse matrix collection were obtained at this meeting.

Sparse Matrix Symposium 1982. W. G. Poole, Jr. was a member of the Advisory Committee for this symposium. A. M. Erisman, R. G. Grimes, J. G. Lewis, W. G. Poole, Jr., and Horst D. Simon attended and presented papers.

Sparse Matrix Methods for Petroleum Engineering. J. G. Lewis, R. G. Grimes and J. L. Phillips have been investigating ordering methods specific to problems from petroleum engineering. The SPARSPAK package has been extended to include an automatic D4 ordering developed by Lewis, and comparisons with standard orderings have been made and presented.

Methods for Solving Symmetric Sparse Linear Equations Out-of-memory. Extensions of the orderings found in SPARSPAK to solve extremely large problems have been developed. Software for these extensions is being prepared. All of the orderings in SPARSPAK, except the Minimum Degree Ordering, have been so extended. A special version of some of this software will be developed shortly, under contract to a major petroleum company, to solve a particular class of linear systems of order up to 250,000 equations.

Preconditioned Conjugate Gradient and Lanczos Methods for Solving Symmetric Linear Systems. H. D. Simon is continuing the research of his thesis topic, and has developed new methods for finding preconditioning matrices for irregular symmetric matrices, particularily those which arise in structural engineering.

Preconditioned Conjugate Gradient Methods for Unsymmetric Linear Systems. E. L. Yip and D. P. Young have separately developed methods for solving unsymmetric sparse linear equations arising in aerodynamic applications, especially in potential flow theory. Dr. Yip's approach has been presented and submitted for publication.

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFOSR-TR- 83-0819	2. GOVT ACCESSION NO. A133657	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Ordering Methods for Sparse Matrices		5. TYPE OF REPORT & PERIOD COVERED Final July 1981 - July 1983
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) John G. Lewis		8. CONTRACT OR GRANT NUMBER(s) Contract F49620-81-C-0072
9. PERFORMING ORGANIZATION NAME AND ADDRESS Energy Technology Applications Division Boeing Computer Services Company, MS 9C-01 565 Andover Park West Tukwila, Washington 98188		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS PE61102F; 2304/A3
11. CONTROLLING OFFICE NAME AND ADDRESS Air Force Office of Scientific Research /NM Building 410 Bolling Air Force Base, D.C. 20332		12. REPORT DATE July 28, 1983
		13. NUMBER OF PAGES 9
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Sparse matrices, reordering algorithms, Hellerman-Rarick, P4, unsymmetric matrices, Gauss elimination

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report summarizes the activities at Boeing Computer Services Company on AFOSR Contract F49620-81-0072 from July 1, 1981 until July 28, 1983. Five tasks were performed: creation of a comprehensive test matrix collection, analysis of the Hellerman-Rarick P4 algorithm, production of a P4 code, production of a diagnostic code, and comparative analysis of several algorithms using the test matrices and the diagnostic code. Reports on the completion of the five tasks are given, relevant reports and publications of project personnel are listed and related sparse matrix activities are discussed.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

END

FILMED

10-83

DTIC